arXiv:1402.4963v3 [cs.CV] 4 Mar 2014

Vesselness via Multiple Scale Orientation Scores

Julius Hannink, Remco Duits and Erik Bekkers

Eindhoven University of Technology, Department of Biomedical Engineering and Department of Mathematics and Computer Science, Eindhoven, the Netherlands
J.Hannink@gmx.de, R.Duits@tue.nl and E.J.Bekkers@tue.nl

Abstract. The multi–scale Frangi vesselness filter is an established tool in (retinal) vascular imaging. However, it cannot properly cope with crossings or bifurcations since it only looks for elongated structures. Therefore, we disentangle crossings/bifurcations via (multiple scale) invertible orientation scores and apply vesselness filters in this domain. This new method via scale–orientation scores performs considerably better at enhancing vessels throughout crossings and bifurcations than the Frangi version. Both methods are evaluated on a public dataset. Performance is measured by comparing ground truth data to the segmentation results obtained by basic thresholding and morphological component analysis of the filtered images.

Keywords: Multi–scale vesselness filters, multi–scale orientation scores, line detection, gauge frames, retinal imaging

1 Introduction

The retinal vasculature enables non–invasive observation of the human circulatory system. A variety of eye–related and systematic diseases such as glaucoma, age–related macular degeneration, diabetes, hypertension, arteriosclerosis or Alzheimer's disease affect the vasculature and may cause functional or geometric changes [1]. Automated quantification of these defects promises massive screenings for systematic and eye–related vascular diseases on the basis of fast and inexpensive imaging modalities, i.e. retinal photography.

To automatically assess the state of the retinal vascular tree, vessel segmentations and/or models have to be created and analyzed. Because retinal images usually suffer from low contrast at small scales, the vasculature needs to be enhanced prior to model creation/segmentation. One well–established approach is the Frangi vesselness filter [2]. It is frequently used in robust retinal vessel segmentation methods [3, 4]. However, the Frangi filter has a known drawback. It cannot properly enhance vessels throughout crossings or bifurcations that make up huge parts of the retinal vascular network.

To generically deal with this issue, we apply the principle of image processing via invertible orientation scores (Fig. 1). In the orientation score domain, crossing/bifurcating lines are disentangled into separate layers corresponding to their orientation (Fig. 2). Consequently, an equivalent of the Frangi filter can be used on this domain to enhance vessels in a crossing–preserving way.

The construction of invertible orientation scores is inspired by the functional architecture of the human cortical columns in the primary visual cortex [5], where decomposition of local orientation allows the visual system to separate crossing/bifurcating structures. Because of this, a human observer robustly identifies the two lines in the exemplary image in Fig. 2, whereas conventional line–filtering on the image domain might fail in cases with noise and low contrast.

Similar approaches of frequency or velocity specific data representations are used in other contexts [6, 7]. The general underlying theme is switching to the corresponding Lie–group of interest in the generic group theoretical approach outlined in [8]. Here, we will develop vesselness filters on the extended Lie–group domains of the rotation translation group $SE(2)$ and the rotation, translation and scaling group $SIM(2)$. Our approach is closely related to the work by Krause, Alles, Burgeth and Weickert who rely on a local Radon transform to disentangle crossings and bifurcations [9].

The extension of vesselness filters on $\mathbb{R}^2$ to $SE(2)$ and $SIM(2)$ allows us to cope with the high degree of (multiple scale) crossings and bifurcations in 2d retinal images. In the end, we show the performance of this new type of vesselness filters by comparison to the multi–scale Frangi vesselness [2], both qualitatively and quantitatively on the High Resolution Fundus (HRF) image dataset available at `http://www5.cs.fau.de/research/data/fundus-images/`.

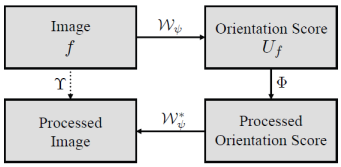

Fig. 1. Image processing via invertible orientation scores.

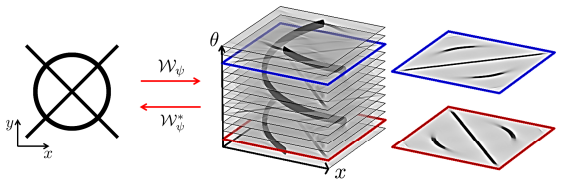

Fig. 2. Exemplary image and corresponding orientation score.

2 Methods

2.1 Orientation Scores on $SE(2)$

An orientation score $U_f : SE(2) \to \mathbb{C}$ is obtained by correlating an input image f with a specially designed, anisotropic wavelet ψ:

$$U_f(\boldsymbol{x},\theta) = \left(\overline{\psi_\theta} \star f\right)(\boldsymbol{x},\theta) = \int_{\mathbb{R}^2} \overline{\psi\left(R_\theta^{-1}(\boldsymbol{y}-\boldsymbol{x})\right)}\, f(\boldsymbol{y})\mathrm{d}\boldsymbol{y} \tag{1}$$

where R_θ denotes a 2d counter–clockwise rotation matrix and $\psi_\theta(\boldsymbol{x}) = \psi(R_\theta^{-1}\boldsymbol{x})$. Exact image reconstruction is achieved by

$$f = \mathcal{F}^{-1}\left[M_\psi^{-1}\mathcal{F}\left[\boldsymbol{x} \mapsto \frac{1}{2\pi}\int_0^{2\pi} (\psi_\theta * U_f)(\boldsymbol{x},\theta)\mathrm{d}\theta\right]\right] \tag{2}$$

where $*$ denotes convolution, $\mathcal{F}[\cdot]$ represents the unitary Fourier transform on $\mathbb{L}_2(\mathbb{R}^2)$ and M_ψ is given by $\int_0^{2\pi} |\mathcal{F}[\psi_\theta]|^2 \mathrm{d}\theta$. Theoretically, reconstruction is well posed for $0 < \delta < M_\psi < \infty$ with arbitrary small δ. Practically, however, it has proven best to aim at $M_\psi \approx 1$ since that ensures optimal stability [10].

One type of wavelets that meet this stability criterion are the cake wavelets described by [11, 12]. They uniformly cover the Fourier domain up to a radius of about the Nyquist frequency ρ_n to satisfy the discrete version of $M_\psi \approx 1$ by design. This is also the general idea behind curvelets and shearlets [13, 14]. They do, however, not give appropriate representations in our context. Curvelets have varying orientation localization over scale, whereas retinal images show detail on all scales. Shearlets are based on the shearing, translation and scaling group, whereas rotation invariance is much more desirable in our context.

The procedure for creating cake wavelets is illustrated in Fig. 3. In detail, the proposed wavelet takes the form

$$\psi(\boldsymbol{x}) = \left(\mathcal{F}^{-1}\left[\boldsymbol{\omega} \mapsto \tilde{\psi}(\rho\cos\varphi, \rho\sin\varphi)\right]\right)(\boldsymbol{x})\; G_s(\boldsymbol{x}) \tag{3}$$

where G_s is an isotropic Gaussian window in the spatial domain and $(\rho, \varphi)^T$ denote polar coordinates in the Fourier domain, i.e. $\boldsymbol{\omega} = (\rho\cos\varphi, \rho\sin\varphi)^T$. The Fourier wavelet $\tilde{\psi}(\boldsymbol{\omega}) = A(\varphi)B(\rho)$ is constructed from

$$A(\varphi) = \begin{cases} B^k\left(\frac{(\varphi \bmod 2\pi) - \pi/2}{s_\theta}\right) & \text{if } \rho > 0 \\ \frac{1}{N_\theta} & \text{if } \rho = 0 \end{cases} \quad , \quad B(\rho) = e^{-\left(\frac{\rho}{t}\right)^2} \sum_{i=0}^{N} \frac{(\rho/t)^{2i}}{i!} \tag{4}$$

$B^k(x)$ denotes the $k-$th order B–spline, N_θ is the number of samples in the orientation direction and $s_\theta = 2\pi/N_\theta$ is the angular stepsize. The function $B(\rho)$ is a Gaussian multiplied with the Taylor series of its inverse up to order N to enforce faster decay. The parameter t is given by $t^2 = 2\hat{\rho}/(1+2N)$ with the inflection point $\hat{\rho}$ that determines the bending point of $B(\rho)$.

As depicted in Fig. 3, the real part of the kernel picks up lines, whereas the imaginary part responds to edges. In the context of vessel filtering, the real part is of primary interest, whereas vessel tracking in the score additionally makes use of the imaginary part [12].

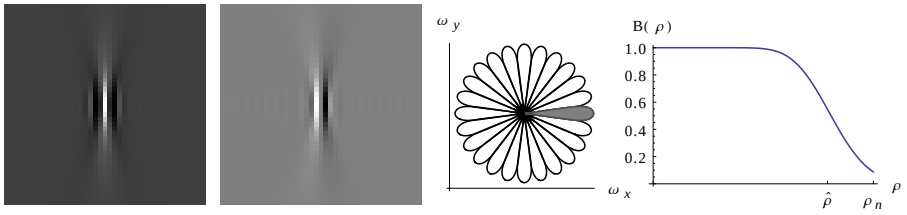

Fig. 3. Real and imaginary part of the cake kernel (left, zoomed), Fourier contours at 70% of the maximum for all orientations and $B(\rho)$ with $\hat{\rho} = 0.8\rho_n$, the Nyquist frequency ρ_n and $N = 8$.

2.2 Gaussian Derivatives in Orientation Scores

The orientation score domain is essentially the 2d special Euclidean motion group $SE(2) \equiv \mathbb{R}^2 \rtimes S^1$ [11, ch.2]. Because of this, all operations Φ on this domain have to be left–invariant to produce a Euclidean invariant net operator Υ on the image [11, ch.2]. This is desirable since the result should be independent on rotation and translation of the input. Φ is left–invariant iff $\Phi \circ \mathcal{L}_{\boldsymbol{g}} = \mathcal{L}_{\boldsymbol{g}} \circ \Phi$ for all $\boldsymbol{g} = (\boldsymbol{x}, \theta) \in SE(2)$ with the shift–twist operator on the score given by

$$\mathcal{L}_{\boldsymbol{g}} U_f(\boldsymbol{h}) = U_f(\boldsymbol{g}^{-1}\boldsymbol{h}) = U_f(R_\theta^{-1}(\boldsymbol{x}' - \boldsymbol{x}), \theta' - \theta), \tag{5}$$

for all $\boldsymbol{g} = (\boldsymbol{x}, \theta), \boldsymbol{h} = (\boldsymbol{x}', \theta') \in SE(2)$. Therefore, we must rely on left–invariant derivatives, given by $\partial_\theta, \partial_\xi = \cos\theta\partial_x + \sin\theta\partial_y$ and $\partial_\eta = -\sin\theta\partial_x + \cos\theta\partial_y$ when constructing vesselness filters on $SE(2)$. These derivatives provide a moving frame of reference on the group steered by the orientation of the wavelet. Their non–zero commutators are given by $[\partial_\theta, \partial_\xi] = \partial_\eta$ and $[\partial_\theta, \partial_\eta] = -\partial_\xi$. Later we will adapt this frame locally to the score, following the theory of best exponential curve fits presented in [11, ch.6]. This compensates for the fact that our wavelet kernel is not always perfectly aligned with all local orientations present in the image (For details, see [11, ch.6]).

Since orientation and spatial direction have different physical units, a conversion factor is needed. This parameter β has unit 1/length and determines the shape of $SE(2)$ geodesics (see Fig. 4). Mathematically, β appears as the only free parameter in the (sub-)Riemannian metric on $SE(2)$ given by

$$d(\boldsymbol{g}_1, \boldsymbol{g}_2) = \inf_{\substack{\gamma(0)=\boldsymbol{g}_1 \\ \gamma(l)=\boldsymbol{g}_2 \\ \dot{\gamma}\in\Delta, l\geq 0}} \int_0^l \sqrt{\mathcal{G}_\beta\big|_{\gamma(s)}\big(\dot{\gamma}(s), \dot{\gamma}(s)\big)}\, \mathrm{d}s \ , \tag{6}$$

with $\Delta = \mathrm{span}\{\partial_\xi, \partial_\eta, \partial_\theta\}$, $\gamma(s) = (\boldsymbol{x}(s), \theta(s))$ and $\mathcal{G}_\beta|_\gamma(\dot{\gamma}, \dot{\gamma}) = \beta^2(\dot{x}\cos\theta + \dot{y}\sin\theta)^2 + \beta^2(-\dot{x}\sin\theta + \dot{y}\cos\theta)^2 + \dot{\theta}^2$ in the Riemannian case. In the sub–Riemannian case, the allowed part of the tangent space is $\Delta = \mathrm{span}\{\partial_\xi, \partial_\theta\}$. The functional in (6) then reduces to $\int_0^l \sqrt{\kappa^2(s) + \beta^2}\, \mathrm{d}s$ with the curvature κ of the spatially projected curve $\boldsymbol{x}(s) = \mathcal{P}_{\mathbb{R}^2}\gamma(s)$ under the condition that $\boldsymbol{g}_2$ is chosen "aligned enough" with $\boldsymbol{g}_1$ (For details, see [15]). Typically, $\beta < 1$.

In order to extract local features in $SE(2)$ with well–posed, left–invariant derivative operators, some regularization has to be included. The only left–invariant diffusion regularization in $SE(2)$ that preserves the non–commutative group structure via the commutators is elliptic diffusion, isotropic w.r.t. the β-metric (6). In this case, regularization is achieved via a spatially isotropic Gaussian with scale $\frac{1}{2}\sigma_s^2$ and a 1d–Gaussian in θ with scale $\frac{1}{2}(\beta\sigma_s)^2$ [11, ch.5]. The regularized derivative operators are convolutions with correspondingly differentiated β-isotropic Gaussians and generalize the concept of Gaussian derivatives used in the Frangi vesselness filter [2] to $SE(2)$.

2.3 Scale–Orientation Scores

To make the kernels described above scale–selective, the pieces of cake have to be further divided. By cutting out pieces in the log–radial direction, they are

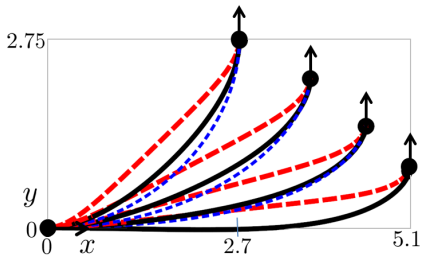

Fig. 4. Spatial projection of $SE(2)$ geodesics for $\beta = 3$ (dashed, red), $\beta = 1$ (solid, black) and $\beta = 1/3$ (thin, dashed, blue). For details on sub–Riemannian geodesics and association fields, see [15].

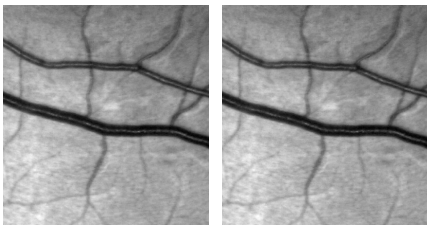

Fig. 5. Exemplary retinal image (left) and reconstruction from its scale–orientation score (right).

made sensitive to a specific frequency range that can be identified with a scale a in the spatial domain. To construct scale selective cake kernels (anisotropic wavelets), [16] uses a radial envelope function

$$B^{\mathrm{MS}}(\rho) = \sum_{l=0}^{N_\rho - 1} B_l^k(\rho) := \sum_{l=0}^{N_\rho - 1} B^k\left(\frac{\log(\rho a^-)}{s_\rho} + l\right) \tag{7}$$

where $B^k(x)$ is the k-th order B–spline function, N_ρ is the total number of scales to sample in the Fourier domain and $s_\rho > 0$ denotes the stepsize in log–scale. The multiplicative character of the spatial scales $a_l = a^- e^{l s_\rho}$ reflects the typical scale transition at bifurcations. Because of the B–spline approach, the scale selective envelopes $B_l^k(\rho)$ sum to one and the $M_\psi \approx 1$ requirement is still met (Fig.6). Scale layers outside a spatially defined range of interest are merged to reduce computational load (Fig. 6). We propose the following multi–scale cake kernel

$$\psi^{\mathrm{MS}}(\boldsymbol{x}) = \left(\mathcal{F}^{-1}\left[M^{-1}\mathcal{F}\left[\tilde{\psi}^{\mathrm{MS}}\right](\boldsymbol{\omega})\right]\right)(\boldsymbol{x}) \tag{8}$$

where $\tilde{\psi}^{\mathrm{MS}}(\boldsymbol{x})$ denotes the wavelet

$$\tilde{\psi}^{\mathrm{MS}}(\boldsymbol{x}) = \left(\mathcal{F}^{-1}\left[\boldsymbol{\omega} \mapsto A(\varphi)B_0^k(\rho)\right]\right)(\boldsymbol{x})\, G_{s_x,s_y}(\boldsymbol{x}) \tag{9}$$

at the finest scale a^-. The anisotropic Gaussian window $G_{s_x,s_y}(\boldsymbol{x})$ reduces long tails along the orientation of the wavelet and suppresses oscillations perpendicular to it induced by narrow sampling bandwidths in $B^{\mathrm{MS}}(\rho)$. The corresponding changes in the Fourier domain are resolved by subsequent normalization with $M(\omega) = N_\rho^{-1} N_\theta^{-1} \sum_{i=1}^{N_\rho} \sum_{j=1}^{N_\theta} a_i^{-1} |\mathcal{F}[\tilde{\psi}^{MS}](a_i R_{\theta_j}^{-1}\boldsymbol{\omega})|$. Thereby, reconstruction is done by summation over scales and angles and perfect invertibility (Fig. 5) is ensured. The scale and orientation specific data representation given by $(\mathcal{W}_\psi f)(\boldsymbol{x}, \theta, a) = (\overline{\psi_\theta^a} \star f)(\boldsymbol{x})$ with $\psi_\theta^a(\boldsymbol{x}) = a^{-1}\psi^{\mathrm{MS}}(a^{-1}R_\theta^{-1}\boldsymbol{x})$ is now set and processing can begin.

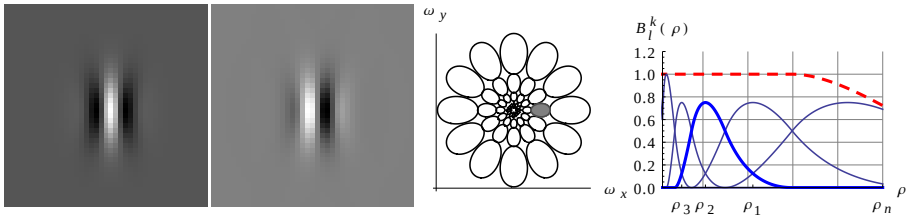

Fig. 6. Real and imaginary part of the multi–scale cake kernel at a_2 (left), Fourier contours of all wavelets (at 70% of maximum) and $B^{\mathrm{MS}}(\rho)$ (red, dashed) with components $B_l^k(\rho)$ (blue). Highlighted components correspond to the kernels shown.

2.4 Vesselness Filtering on Scale–Orientation Scores

The single–scale Frangi vesselness filter $\mathcal{V}_0^{\mathrm{Fr}}$ makes use of an anisotropy measure $\mathcal{R}$ and a structure measure $\mathcal{S}$ based on second order derivatives in a coordinate system aligned with elongated structures [2]. This approach is now generalized to (scale–)orientation scores for crossing–preserving vesselness filtering. Given a convexity criterion $\mathcal{Q} > 0$ on transversal vessel profiles and the measures $\mathcal{R}$ resp. $\mathcal{S}$, the $SE(2)$ vesselness expression $\mathcal{V}_0^{SE(2)}(\mathcal{U}_f^a) : SE(2) \to \mathbb{R}^+$ is identical to the one proposed by [2]:

$$\mathcal{V}_0^{SE(2)}\left(\mathcal{U}_f^a\right) = \begin{cases} 0 & \text{if } \mathcal{Q} \leq 0 \\ \exp\left(-\frac{\mathcal{R}^2}{2\sigma_1^2}\right)\left[1 - \exp\left(-\frac{\mathcal{S}}{2\sigma_2^2}\right)\right] & \text{if } \mathcal{Q} > 0 \end{cases} \tag{10}$$

where $\mathcal{U}_f^a(\boldsymbol{x}, \theta) = (\mathcal{W}_\psi f)(\boldsymbol{x}, \theta, a)$, $a > 0$ fixed, is a single scale layer of the multi–scale wavelet transform. Typically, $\sigma_1 = 0.5$ and $\sigma_2 = 0.2\,||\mathcal{S}||_\infty$.

There are two natural generalizations of $\mathcal{V}_0^{\mathrm{Fr}}$ to $SE(2)$ that differ in the choice of coordinate system used to define $\mathcal{R}, \mathcal{S}$ and $\mathcal{Q}$. One option is to work in the moving frame of reference $\{\partial_\xi, \partial_\eta, \partial_\theta\}$. Since ξ and η are respectively parallel and orthogonal to the orientation in each θ–slice, second order Gaussian derivatives along these directions carry the same information as the eigenvalues of the Hessian used in $\mathcal{V}_0^{\mathrm{Fr}}$ [2]. Therefore, $\mathcal{R}, \mathcal{S}$ and $\mathcal{Q}$ are computed as

$$\mathcal{R} = \frac{(\partial_\xi^2 \mathcal{U}_f^a)^{s,\beta}}{(\partial_\eta^2 \mathcal{U}_f^a)^{s,\beta}} \quad , \quad \mathcal{S} = \left[(\partial_\xi^2 \mathcal{U}_f^a)^{s,\beta}\right]^2 + \left[(\partial_\eta^2 \mathcal{U}_f^a)^{s,\beta}\right]^2 \quad , \quad \mathcal{Q} = (\partial_\eta^2 \mathcal{U}_f^a)^{s,\beta} \tag{11}$$

where the superscripts $^{s,\beta}$ indicate Gaussian derivatives at spatial scale $s = \frac{1}{2}\sigma_s^2$ and angular scale $\frac{1}{2}(\beta\sigma_s)^2$. The generalization of the filter in the $\{\partial_\xi, \partial_\eta, \partial_\theta\}$ frame is referred to as $\mathcal{V}_0^{\xi,\eta}$ at single scales and as $\mathcal{V}^{\xi,\eta}$ in the multiple scale recombination, similar to the notation in [2]. The other possible coordinate system is the Gauge frame $\{\partial_{\boldsymbol{a}}, \partial_{\boldsymbol{b}}, \partial_{\boldsymbol{c}}\}$ determined by the eigendirections of the Hessian $(\mathcal{H}^{s,\beta}\mathcal{U}_f^a)(\boldsymbol{g})$ at scale s and $\boldsymbol{g} \in SE(2)$, normalized w.r.t the β–metric (6). As the filter is no longer confined to θ-slices because the Gauge frame is free to fully align with the data in the score, the analogies to [2] are even stronger in this

frame. Given the eigenvalues of the Hessian $(\mathcal{H}^{s,\beta}\mathcal{U}_f^a)(\boldsymbol{g})$ ordered in absolute magnitude $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, $\mathcal{R}, \mathcal{S}$ and $\mathcal{Q}$ are computed as

$$\mathcal{R} = \frac{\lambda_1}{c} \quad , \quad \mathcal{S} = \lambda_1^2 + c^2 \quad , \quad \mathcal{Q} = c \tag{12}$$

with $c = \frac{1}{2}(\lambda_2 + \lambda_3)$. As such, c is comparable to the orientation confidence defined by [11]. The generalization of the vesselness filter in this frame is referred to as $\mathcal{V}_0^{\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}}$ at single scales and as $\mathcal{V}^{\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}}$ in the multi–scale recombination.

The generalization of the multi–scale Frangi filter $\mathcal{V}^{\mathrm{Fr}}$ to the $SIM(2)$ domain is achieved by image reconstruction from vesselness filtered scale–orientation scores with subsequent intensity normalization

$$\big(\mathcal{V}^{SIM(2)}(f)\big)(\boldsymbol{x}) = \mu_\infty^{-1} \sum_{i=1}^{N_s} \mu_{i,\infty}^{-1} \sum_{j=1}^{N_\theta} \big(\mathcal{V}_0^{SE(2)}(U_f^{a_i})\big)(\boldsymbol{x}, \theta_j) \tag{13}$$

where μ_∞ and $\mu_{i,\infty}$ are the maximum values, i.e. $||\cdot||_\infty$ norms, taken over the subsequent sums. To get comparable results, the multi–scale Frangi vesselness filter is also computed via summation over single scale results and normalization by the maximum value.

Fig. 7 shows multi–scale vesselness results on an exemplary retinal image obtained with the Frangi filter and our two methods for five scales $\{1.5, 2.4, 3.8, 6.0, 9.5\}$ px, $\beta = 0.05$ and 12 orientations sampled in $[0, \pi)$. Both our methods clearly outperform the Frangi filter at crossings and bifurcations. The Gauge–frame method $\mathcal{V}^{\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}}$ gives best results since it is better aligned with elongated structures in the score.

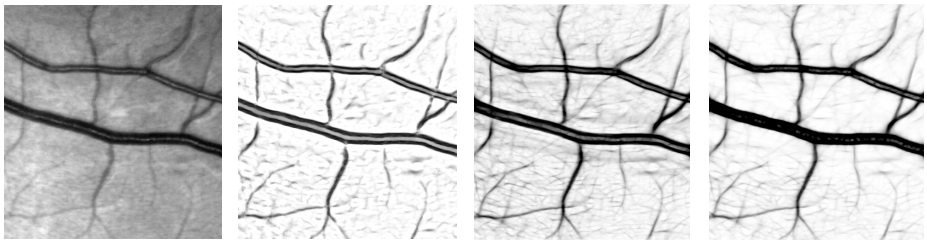

Fig. 7. Retinal image f and multi–scale vesselness filtering results for the Frangi filter $\mathcal{V}^{\mathrm{Fr}}(f)$ and our two methods $\mathcal{V}^{\xi\eta}(f)$ resp. $\mathcal{V}^{\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}}(f)$ (left to right). Gray scales adjusted.

3 Experiments

To show the benefit of crossing–preservation in multiple scale vesselness filtering, we devised a simple segmentation algorithm to turn a vesselness filtered image $\mathcal{V}(f)$ into a binary vessel map. First, an adaptive thresholding is applied, yielding a binary image

$$f_B = \Theta\big([\mathcal{V}(f) - G_\gamma * \mathcal{V}(f)] - t\big) \tag{14}$$

where Θ is the Heaviside step function and G_γ is a Gaussian of scale $\gamma \gg 1$. In a second step, the connected morphological components in f_B are subject to size and elongation constraints. Components counting less than τ pixels or showing elongations below a threshold ν are removed. The parameters γ, τ and ν are fixed at 100 px, 500 px and 0.85 respectively. $\mathcal{V}(f)$ is either obtained via $\mathcal{V}^{\mathrm{Fr}}$ or via the $SIM(2)$ method $\mathcal{V}^{\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}}$ with the settings mentioned earlier.

This segmentation algorithm is evaluated on the HRF dataset consisting of wide–field fundus images for a healthy, diabetic retinopathy and glaucoma group (15 images each, ground truths provided). Average sensitivity and accuracy on the whole dataset are shown in Fig. 8 over threshold values t. Our method via invertible scale–orientation scores performs considerably better than the method based on the multi–scale Frangi filter. The segmentation results obtained with $\mathcal{V}^{\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}}$ are more stable w.r.t variations in the threshold t and the performance on the small vasculature has improved as measured via the sensitivity. Average sensitivity, specificity and accuracy at a threshold $t = 0.05$ resp. given by $0.786, 0.988, 0.969$ (healthy), $0.811, 0.963, 0.953$ (diabetic retinopathy) and $0.641, 0.988, 0.960$ (glaucoma) compare well with other algorithms evaluated on the HRF dataset (see [3, Tab. 5]). On the diabetic retinopathy group, our method even outperforms existing segmentation methods. Fig. 8 shows a full segmentation computed with the proposed method and an in–detail patch.

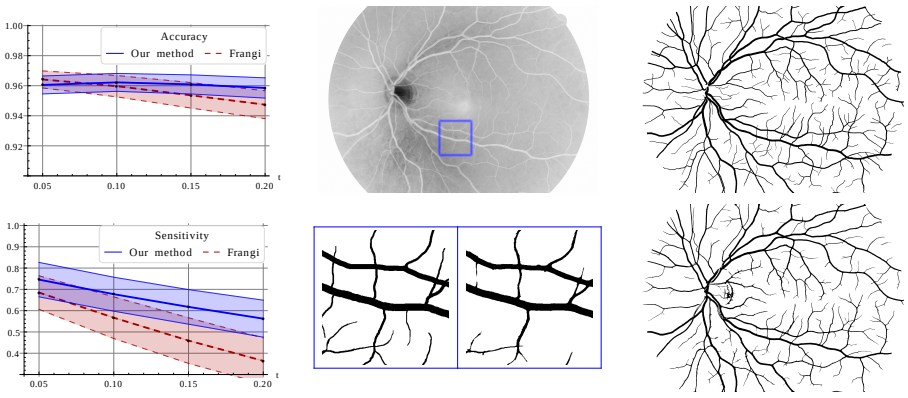

Fig. 8. Average accuracy and sensitivity on the HRF dataset over threshold values t. Shaded regions correspond to $\pm 1\,\sigma$ (left). Retinal image and patch–wise ground truth/segmentation at $t = 0.05$ (center). Full ground truth and segmentation (right).

4 Discussion

We developed (multi–scale) crossing–preserving vesselness filters as generalizations of [2] to the extended Lie group domains $SE(2)$ resp. $SIM(2)$. The new filters were evaluated qualitatively and quantitatively on a public dataset and outperformed the Frangi filter and existing segmentation methods. This shows the method's potential for application in other areas of vascular imaging. Reduced sensitivity on the glaucoma group is most likely due to non–uniform contrast in the input. Future work therefore includes contrast normalization as preprocessing and concatenation with enhancements [11, 16] and tracking [12].

Acknowledgements: The research leading to these results has received funding from the ERC councilunder the EC's 7th Framework Programme (FP7/2007–2013) / ERC grant agr. No. 335555.

References

[1] Ikram, M.K., Ong, Y.T., Cheung, C.Y., Wong, T.Y.: Retinal Vascular Caliber Measurements: Clinical Significance, Current Knowledge and Future Perspectives. Ophthalmologica **229**(3) (2013) 125–136
[2] Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. LNCS **1496** (1998) 130–137
[3] Budai, A., Bock, R., Maier, A., Hornegger, J., Michelson, G.: Robust Vessel Segmentation in Fundus Images. IJBI **2013** (2013)
[4] Lupascu, C.A., Tegolo, D., Trucco, E.: FABC: Retinal Vessel Segmentation Using AdaBoost. IEEE T-ITB **14**(5) (2010) 1267–1274
[5] Duits, R., Felsberg, M., Granlund, G., ter Haar Romeny, B.: Image Analysis and Reconstruction using a Wavelet Transform Constructed from a Reducible Representation of the Euclidean Motion Group. IJCV **72**(1) (2007) 79–102
[6] Duits, R., Janssen, B., Bruurmijn, M., Florack, L., Van Assen, H.: Evolution Equations on Gabor Transforms and their Applications. ACHA (to appear 2014)
[7] Barbieri, D., Citti, G., Cocci, G., Sarti, A.: A cortical–inspired geometry for contour perception and motion integration. arXiv preprint arXiv:1301.3433 (2013)
[8] Duits, R., Burgeth, B.: Scale Spaces on Lie Groups. In Sgallari, Murli, Paragios, eds.: Scale Space and Variational Methods, Springer (2007) 300–312
[9] Krause, M., Alles, R.M., Burgeth, B., Weickert, J.: Fast retinal vessel analysis. JRTIP (2013) 1–10
[10] Fuehr, H.: Abstract Harmonic Analysis of Continuous Wavelet Transforms. Springer (2005)
[11] Franken, E.: Enhancement of Crossing Elongated Structures in Images. PhD thesis, Technical University Eindhoven (2008)
[12] Bekkers, E., Duits, R., Berendschot, T., ter Haar Romeny, B.: A Multi–Orientation Analysis Approach to Retinal Vessel Tracking. JMIV (2014) 1–28
[13] Candes, E., Demanet, L., Donoho, D., Ying, L.: Fast Discrete Curvelet Transforms. SIAM MMS **5**(3) (2006) 861–899
[14] Bodmann, B.G., Kutyniok, G., Zhuang, X.: Gabor Shearlets. arXiv preprint arXiv:1303.6556 (2013)
[15] Duits, R., Boscain, U., Rossi, F., Sachkov, Y.: Association Fields via Cuspless Sub–Riemannian Geodesics in SE (2). JMIV **1** (2013) 32
[16] Sharma, U., Duits, R.: Left-invariant evolutions of wavelet transforms on the Similitude Group. arXiv preprint arXiv:1306.1800 (2013)